Data Center Server Provision: Distributed Asynchronous Control for Coupled Renewal Systems

Xiaohan Wei and Michael J. Neely

arXiv:1512.01609v1 [cs.DC] 5 Dec 2015

***Abstract*—This paper considers a cost minimization problem for data centers with N servers and randomly arriving service requests. A central router decides which server to use for each new request. Each server has three types of states (active, idle, setup) with different costs and time durations. The servers operate asynchronously over their own states and can choose one of multiple sleep modes when idle. We develop an online distributed control algorithm so that each server makes its own decisions and the overall time average cost is near optimal with probability 1. The algorithm does not need probability information for the arrival rate or job sizes. Next, an improved algorithm that uses a single queue is developed via a "virtualization" technique. The improvement is shown to provide the same (near optimal) costs while significantly reducing delay, as shown in simulations.**

I. Introduction

Consider a data center that consists of a central controller and N servers that serve randomly arriving requests. The system operates in slotted time with time slots $t \in \{0, 1, 2, \ldots\}$. Each server $n \in \{1, \ldots, N\}$ has three basic states:

- Active: The server is available to serve requests. Server n incurs a cost of $e_n \geq 0$ on every active slot, regardless of whether or not requests are available to serve.
- Idle: A low cost sleep state where no requests can be served. The idle state is actually comprised of a choice of multiple sleep modes with different per-slot costs. The specific sleep mode also affects the setup time required to transition from the idle state to the active state.
- Setup: A transition period from idle to active during which no requests can be served. The setup cost and duration depend on the preceding sleep mode. The setup duration is typically more than one slot, and can be a random variable that depends on the server n and on the preceding sleep mode.

An active server can choose to transition to the idle state at any time. When it does so, it chooses the specific sleep mode to use and the amount of time to sleep. For example, deeper sleep modes can shut down more electronics and thereby save on per-slot idling costs. However, a deeper sleep incurs a longer setup time when transitioning back to the active state. Each server makes separate decisions about when to transition and what sleep mode to use. The resulting transition times for each server are asynchronous. On top of this, a central controller makes slot-wise decisions for routing requests to servers. It can also reject requests (with a certain amount of cost) if it decides they cannot be supported. The goal is to minimize the overall time average cost.

This problem is challenging mainly for two reasons: First, since each setup state generates cost but serves no request, it is not clear whether or not transitioning to idle from the active state indeed saves power. It is also not clear which sleep mode the server should switch to. Second, if one server is currently in a setup state, it cannot make another decision until it reaches the active state (which typically takes more than one slot), whereas other active servers can make decisions during this time. Thus, this problem can be viewed as a system with coupled Markov decision processes (MDPs) making decisions asynchronously. Classical methods for MDPs, such as dynamic programming and linear programming [1][2][3], can be impractical for two reasons: First, the state space has dimension that depends on the number of servers, making solutions difficult when the number of servers is large. Second, some statistics of the system, such as the arrival probabilities, are unknown.

In this paper, we use renewal-reward theory (e.g. [4]) together with Lyapunov optimization (e.g. [5]) to develop a simple implementable algorithm for this problem. Our work is not alone in approaching the problem this way. Work in [6] uses Lyapunov theory for a system with one renewal server, while work in [7] considers a multi-server system but in a deterministic context. The work in [8] develops a two stage algorithm for stochastic multi-server systems, but the first stage must be solved offline. This paper is distinct in that we overcome the open challenges of previous work and develop a near optimal fully online algorithm for a stochastic system with multiple servers.

Several alternative approaches to multi-server systems use queueing theory. Specifically, the work in [9] treats the multi-server system as an $M/M/k/setup$ queue, whereas the works in [10] and [11] consider modeling a single server as a multi-class $M/G/1$ queue. These approaches require Poisson traffic and do not treat the asynchronous control problem with multiple sleep options considered in the current paper.

Experimental work on power and delay minimization is treated in [12], which proposes to turn each server ON and OFF according to the rule of an $M/M/k/setup$ queue. The work in [13] applies Lyapunov optimization to optimize power in virtualized data centers, which assumes each server has negligible setup time and ON/OFF decisions are made synchronously at each server. The work [14] solves a deterministic

The authors are with the Electrical Engineering department at the University of Southern California, Los Angeles, CA.

optimization problem assuming the future is known. The works [15], [16] focus on power-aware provisioning over a time scale large enough so that the whole data center can adjust its service capacity. Specifically, [15] considers load balancing across geographically distributed data centers, and [16] considers provisioning over a finite time interval and introduces an online 3-approximation algorithm.

To our knowledge, the current paper is the first to consider the stochastic control problem with heterogeneous servers and multiple idle and setup states at each server. This gives rise to a nonstandard asynchronous problem with coupled Markov decision systems. A novel aspect of the solution is the construction of a process with super-martingale properties by piecing together the asynchronous processes at each server. This is interesting because neither the individual server processes nor their asynchronous sum are super-martingales. The technique yields a simple distributed algorithm that can likely be applied more broadly for other coupled stochastic MDPs.

II. Problem formulation

Consider a slotted time system with N servers, denoted by set $\mathcal{N}$, that serve randomly incoming requests.

A. Front-end model

At each time slot $t \in \{0, 1, 2, \ldots\}$, $\lambda(t)$ new requests arrive at the system. We assume $\lambda(t)$ takes values in a finite set Λ. Let $R_n(t),\ n \in \mathcal{N}$ denote the number of requests routed into server n at time t. In addition, the system is allowed to reject requests. Let $d(t)$ be the number of requests that are rejected on slot t, and let $c(t)$ be the corresponding per-request cost for such rejection. Assume $c(t)$ takes values in a finite set $\mathcal{C}$. The $R_n(t)$ and $d(t)$ decision variables on slot t must be nonnegative integers that satisfy:

$$\sum_{n=1}^{N} R_n(t) + d(t) = \lambda(t)$$
$$\sum_{n=1}^{N} R_n(t) \leq R_{\max}$$

for a given integer $R_{max} > 0$. The vector process $(\lambda(t), c(t))$ takes values in $\Lambda \times \mathcal{C}$ and is assumed to be an independent and identically distributed (i.i.d.) vector over slots $t \in \{0, 1, 2, \ldots\}$ with an unknown probability mass function.

Each server n maintains a request queue $Q_n(t)$ that stores the requests that are routed to it. Requests are served in a FIFO manner with queueing dynamics as follows:

$$Q_n(t+1) = \max\left\{Q_n(t) + R_n(t) - \mu_n(t)H_n(t),\ 0\right\}. \quad (1)$$

where $H_n(t)$ is an indicator variable that is 1 if server n is active on slot t, and 0 else, and $\mu_n(t)$ is a random variable that represents the number of requests can be served on slot t. Each queue is initialized to $Q_n(0) = 0$. Assume that, every slot in which server n is active, $\mu_n(t)$ is independent and identically distributed with a known mean μ_n, and upper bounded with $\mu(t) \leq \mu_{\max},\ \forall t$. This randomness can model variation in job sizes.

Assumption 1: The process $\{(\lambda(t), c(t))\}_{t=0}^{\infty}$ is observable, i.e. the router can observe the $(\lambda(t), c(t))$ realization each time slot t before making decisions. In contrast, the process $\{\mu_n(t)\}_{t=0}^{\infty}$ is not observable, i.e. given that $H_n(t) = 1$, at the beginning of slot t server n knows a random service will take place, but it does not know the realization of $\mu_n(t)$ until the end of slot t.

B. Back-end server model

Each server $n \in \mathcal{N}$ has three types of states: active, idle, and setup. The idle state of each server n is further decomposed into a collection of distinct idle modes. Each server $n \in \mathcal{N}$ makes decisions over its own *renewal frames*. Define the renewal frame for server n as the time period between successive visits to active state (with each renewal period ending in an active state). Let $T_n[f]$ denote the frame size of the fth renewal frame for server n, for $f \in \{0, 1, 2, \ldots\}$. Let $t_f^{(n)}$ denote the start of frame f, so that $T_n[f] = t_{f+1}^{(n)} - t_f^{(n)}$. Assume that $t_0^{(n)} = 0$ for all $n \in \mathcal{N}$, so that time slot 0 the start of the first renewal frame (labeled frame $f = 0$) for all servers. For simplicity, assume all servers are "active" on slot $t = -1$. Thus, the slot just before each renewal frame is an active slot.

Fix a server $n \in \mathcal{N}$ and a frame index $f \in \{0, 1, 2, \ldots\}$. Time $t_f^{(n)}$ marks the start of renewal frame f. At this time, server n must decide whether to remain active or to go idle. If it remains active then the renewal frame lasts for one slot, so that $T_n[f] = 0$. If it goes idle, it chooses an idle mode from a finite set $\mathcal{L}_n$, representing the set of idle mode options. Let $\alpha_n[f]$ represent this initial decision for server n at the start of frame f, so that:

$$\alpha_n[f] \in \{active\} \cup \mathcal{L}_n$$

where $\alpha_n[f] = active$ means the server chooses to remain active. If the server chooses to go idle, so that $\alpha_n[f] \in \mathcal{L}_n$, it then chooses a variable $I_n[f]$ that represents *how much time it remains idle*. The decision variable $I_n[f]$ is chosen as an integer in the set $\{1, \ldots, I_{max}\}$ for some given integer $I_{max} > 0$. The consequences of these decisions are described below.

- Case $\alpha_n[f] = active$. The frame starts at time $t_f^{(n)}$ and has size $T_n[f] = 1$. The active variable becomes $H_n(t_f^{(n)}) = 1$ and an activation cost of e_n is incurred on this slot $t_f^{(n)}$. A random service variable $\mu_n(t_f^{(n)})$ is generated and requests are served according to the queue update (1). Recall that, under Assumption 1, the value of $\mu_n(t)$ is not known until the end of the slot.
- Case $\alpha_n[f] \in \mathcal{L}_n$. In this case, the server chooses to go idle and $\alpha_n[f]$ represents the specific sleep mode chosen. The idle duration $I_n[f]$ is also chosen as an integer in the set $[1, I_{max}]$. After the idle duration completes, the setup duration starts and has an independent and random duration $\tau_n[f] = \hat{\tau}(\alpha_n[f])$, where $\hat{\tau}(\alpha_n[f])$ is an integer random variable with a known mean and variance that depends on the sleep mode $\alpha_n[f]$. At the end of the setup time the system goes active and serves with a random $\mu_n(t)$ as before. The active variable is $H_n(t) = 0$ for all

TABLE I
CONTROL PARAMETERS

parameters	control objective
$R_n(t)$	requests routed to server n at slot t
$d(t)$	requests rejected at slot t
$\alpha_n[f]$	the option server n takes at frame f
$I_n[f]$	number of slots server n stays idle at frame f

slots t in the idle and setup times, and is 1 at the very last slot of the frame. Further:

- Idle cost: Every slot t of the idle time, an idle cost of $g_n(t) = \hat{g}_n(\alpha_n[f])$ is incurred (so that the idle cost depends on the sleep mode). The process $g_n(\tau) = 0$ if server n is not idle on slot τ. The idle cost can be zero, but can also be a small but positive value if some electronics are still running in the sleep mode chosen.
- Setup cost: Every slot t of the setup time, a cost of $W_n(t) = \hat{W}_n(\alpha_n[f])$ is incurred. The process $W_n(\tau) = 0$ if server n is not in a setup duration on slot τ.

Thus, the length of frame f for server n is:

$$T_n[f] = \begin{cases} 1, & \text{if } \alpha_n[f] = active; \\ I_n[f] + \tau_n[f] + 1, & \text{if } \alpha_n[f] \in \mathcal{L}_n. \end{cases} \quad (2)$$

In summary, the costs $\hat{g}_n(\alpha_n)$, $\hat{W}_n(\alpha_n)$ are *deterministic functions* of $\alpha_n \in \mathcal{L}_n$ whereas the setup time $\hat{\tau}_n(\alpha_n)$ is a *random function* of $\alpha_n \in \mathcal{L}_n$. We further make the following assumption regarding $\hat{\tau}_n(\alpha_n)$:

Assumption 2: For any $\alpha_n \in \mathcal{L}_n$, the function $\hat{\tau}_n(\alpha_n)$ is an integer random variable with known mean and variance, as well as bounded first four moments. Denote $\mathbb{E}[\tau_n(\alpha_n)] = m_{\alpha_n}$ and $\text{Var}[\tau_n(\alpha_n)] = \sigma^2_{\alpha_n}$.

Table I summarizes the control variables introduced in this section. The data center architecture is shown is Fig. 1.

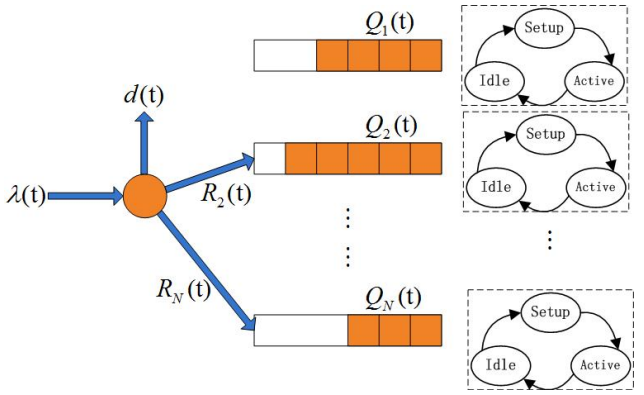

Fig. 1. Illustration of N servers system with N request queues.

C. Performance objective

For each $n \in \mathcal{N}$, let $\overline{C}$, $\overline{W}_n$, $\overline{E}_n$, $\overline{G}_n$ be the time average costs resulting from rejection, setup, service and idle, respectively. They are defined as follows: $\overline{C} = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} d(t)c(t)$, $\overline{W}_n = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} W_n(t)$, $\overline{E}_n = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} e_n H_n(t)$, $\overline{G}_n = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} g_n(t)$.

The goal is to design a joint routing and service policy so that the time average overall cost is minimized and all queues are stable, i.e.

$$\min \ \overline{C} + \sum_{n=1}^{N} \left(\overline{W}_n + \overline{E}_n + \overline{G}_n\right), \text{ s.t. } Q_n(t) \text{ stable } \forall n. \quad (3)$$

The constraint in (3) is not easy to work with. In order to get an optimization problem one can deal with, we further define the time average request rate, rejection rate, routing rate and service rate as $\overline{\lambda}$, $\overline{d}$, $\overline{R}_n$, and $\overline{\mu}_n$ respectively: $\overline{\lambda} = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \lambda(t)$, $\overline{d} = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} d(t)$, $\overline{R}_n = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} R_n(t)$, $\overline{\mu}_n = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \mu_n(t) H_n(t)$.

Then, rewrite the problem (3) as follows

$$\min \quad \overline{C} + \sum_{n=1}^{N} \left(\overline{W}_n + \overline{E}_n + \overline{G}_n\right) \quad (4)$$

$$\text{s.t.} \quad \overline{R}_n \le \overline{\mu}_n, \ \forall n \in \mathcal{N} \quad (5)$$

$$\sum_{n=1}^{N} R_n(t) \le R_{\max}, \ \sum_{n=1}^{N} R_n(t) + d(t) = \lambda(t) \ \forall t \quad (6)$$

Constraint (5) and (6) ensure the service rate at queue n is large enough to support its arrival rate. We aim to develop an algorithm so that each server can make its own decision (without looking at the workload or service decision of any other server) and prove its near optimality.

III. AN ONLINE CONTROL ALGORITHM

We now present an online control algorithm which makes routing and service decisions as functions of $(\lambda(t), c(t), \mathbf{Q}(t))$, where $\mathbf{Q}(t) = (Q_1(t), \ldots, Q_N(t))$. The algorithm is inherited from a classical Lyapunov optimization technique and features a trade-off parameter $V > 0$. It does not require knowledge about the statistics of $\lambda(t)$, $c(t)$ and yields a simple distributed implementation. Specifically, each server makes frame-based decisions regarding its own queue state $Q_n(t)$, whereas the system makes routing and rejection decisions slot-wise based on $(\lambda(t), c(t), \mathbf{Q}(t))$.

A. Thresholding algorithm on admission and routing

During each time slot, the data center chooses nonnegative integers $R_n(t)$ and $d(t)$ to solve the following optimization problem:

$$\min \quad Vd(t)c(t) + \sum_{n=1}^{N} Q_n(t) R_n(t) \quad (7)$$

$$\text{s.t.} \quad \sum_{n=1}^{N} R_n(t) + d(t) = \lambda(t), \ 0 \le \sum_{n=1}^{N} R_n(t) \le R_{\max}.$$

Notice that the solution of this problem admits a simple thresholding rule, i.e.

$$d(t) = \begin{cases} \max\{\lambda(t) - R_{\max}, \ 0\}, & \text{if } \exists n \in \mathcal{N} \text{ s.t.} \\ & Q_n(t) \le Vc(t); \\ \lambda(t), & \text{otherwise.} \end{cases} \quad (8)$$

$$R_n(t) = \begin{cases} \min\{\lambda(t), R_{\max}\}, & \text{if } Q_n(t) \text{ is the shortest} \\ & \text{queue and } Q_n(t) \leq Vc(t); \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

B. Ratio of expectation minimization on service decisions

Each server makes its own frame-based service decisions. Specifically, for server n, at the beginning of its f-th renewal frame $t_f^{(n)}$, it observes its current queue state $Q(t_f^{(n)})$ and makes decisions on $\alpha_n[f] \in \{active\} \cup \mathcal{L}_n$ and $I_n[f]$ so as to solve (11) (see single column equations in a later page for (11)-(14)), where $B_0 = \frac{1}{2}(R_{\max}+\mu_{\max})\mu_{\max}$. Recall the definition of $T_n[f]$, if the server chooses to remain active, then the frame length is exactly 1, otherwise, the server is allowed to choose how long it stays in idle. It can be easily shown that over all randomized decisions between staying active and going to different idle states, it is optimal to make a pure decision which either stays active or goes to one of the idle states with probability 1. Thus, we are able to simplify the problem by computing $D_n[f]$ for active and idle options separately.

- If the server chooses to go active, i.e. $\alpha_n[f] = active$, then,
$$D_n[f] = Ve_n - Q_n(t_f^{(n)})\mu_n. \tag{10}$$
- If the server chooses to go idle, i.e. $\alpha_n[f] \in \mathcal{L}_n$, then, $D_n[f]$ is shown in (12), which follows from the fact that if the server goes idle, then, $H_n(t)$ are all zero during the frame except for the last slot. Now we compute the optimal idle option $\alpha_n[f] \in \mathcal{L}_n$ and idle time length $I_n[f]$ given the server chooses to go idle. The following lemma illustrates that the decision on $I_n[f]$ can also be reduced to pure decisions.

  Lemma 1: The best decision minimizing (12) is a pure decision which takes one $\alpha_n[f] \in \mathcal{L}_n$ and one integer value $I_n[f] \in \{1, \cdots, I_{\max}\}$ minimizing the deterministic function (13).

  The proof of above lemma is given in appendix A.

Then, the server computes the minimum of (13), which goes in the following two steps:

1) For each $\alpha_n \in \mathcal{L}_n$, compute the optimal $I_n[f]$, which is one of the two integers closest to (14) which achieves the smaller value on (13), where $\lfloor\cdot\rfloor$ and$\lceil\cdot\rceil$ stand for floor function and ceiling function respectively.
2) Compare (13) for different $\alpha_n \in \mathcal{L}_n$ and choose the one achieving the minimum.

Then, the server compares (10) with the minimum of (13). If (10) is less than the minimum of (13), then, the server chooses to go active. Otherwise, the server chooses to go idle and stays idle for the time length given the one achieving the minimum of (13).

Overall, our algorithm is summarized in Algorithm 1.

IV. Performance Analysis

In this section, we prove the online algorithm introduced in the last section makes all request delay queues $Q_n(t)$ bounded (on the order of V) and achieves the near optimality with sub optimality gap on the order of $1/V$.

Algorithm 1:

- At each time slot t, the data center observes $\lambda(t)$, $c(t)$, and $\mathbf{Q}(t)$. It chooses rejection variable $d(t)$ according to (8) and routing variable $R_n(t)$ according to (9).
- For each server $n \in \mathcal{N}$, at the beginning of its f-th frame $t_f^{(n)}$, observe its queue state $Q_n(t_f^{(n)})$ and compute (10) and the minimum of (13). If (10) is less than the minimum of (13), then the server still stays active. Otherwise, the server switches to the idle state minimizing (13) and stays idle for $I_n[f]$ achieving the minimum of (13).
- Update $Q_n(t),\ \forall n \in \mathcal{N}$ according to (1).

A. Bounded request queues

Lemma 2: If $Q_n(0) = 0,\ \forall n \in \mathcal{N}$, then, each request queue $Q_n(t)$ is deterministically bounded with bound: $Q_n(t) \leq Vc_{\max} + R_{\max},\ \forall t,\ \forall n \in \mathcal{N}$, where $c_{\max} \triangleq \max_{c\in\mathcal{C}} c$.

Proof: We use induction to prove the claim. Base case is trivial since $Q_n(0) = 0 \leq Vc_{\max} + R_{\max}$. Suppose the claim holds at the beginning of $t = i$ for $i > 0$, so that $Q_n(i) \leq Vc_{\max} + R_{\max}$. Then,

1) If $Q_n(i) \leq Vc_{\max}$, then, it is possible for the queue to increase during slot i. However, the increase of the queue within one slot is bounded by $R_{\max}$. This implies at the beginning of slot $i+1$, $Q_n(i+1) \leq Vc_{\max} + R_{\max}$.
2) If $Vc_{\max} < Q_n(i) \leq Vc_{\max} + R_{\max}$, then, according to (9), it is impossible to route any request to server n during slot i, and $R_n(i) = 0$ which results in $Q_n(i+1) \leq Vc_{\max} + R_{\max}$.

This completes the proof. ■

Lemma 3: The proposed algorithm meets the constraint (5) with probability 1.

Proof: From the queue update rule (1), it follows, $Q_n(t+1) \geq Q_n(t) + R_n(t) - \mu_n H_n(t)$. Taking telescoping sums from 0 to $T-1$ gives $Q_n(T) \geq Q_n(0) + \sum_{t=0}^{T-1} R_n(t) - \sum_{t=0}^{T-1} \mu_n H_n(t)$. Since $Q_n(0) = 0$, dividing both sides by T gives $\frac{Q_n(T)}{T} \geq \frac{1}{T}\sum_{t=0}^{T-1} R_n(t) - \frac{1}{T}\sum_{t=0}^{T-1} \mu_n H_n(t)$. Substituting the bound $Q_n(T) \leq Vc_{\max} + R_{\max}$ from lemma 2 into the above inequality and taking limit as $T \to \infty$ give the desired result. ■

B. Optimal randomized stationary policy

In this section, we introduce a class of algorithms which are theoretically helpful for doing analysis, but practically impossible to implement.

Since servers are coupled only through time average constraint (5), each server n can be viewed as a separate renewal system, thus, it can be shown that any possible time average service rate $\overline{\mu}_n$ can be achieved through a frame based stationary randomized service decision, meaning that the decisions are i.i.d. over frames (see [6] for details). Furthermore, it can be shown that the optimality of (4)-(6) can be achieved over the following randomized stationary algorithms: At the beginning of each time slot t, the data center observes the incoming requests $\lambda(t)$ and rejecting cost $c(t)$, then routes

$$D_n[f] \triangleq \frac{\mathbb{E}\left[\sum_{t=t_f^{(n)}}^{t=t_{f+1}^{(n)}-1}\left(VW_n(t)+Ve_nH_n(t)+Vg_n(t)-Q_n(t_f^{(n)})\mu_n(t)H_n(t)\right)+\left(t-t_f^{(n)}\right)B_0 \;\middle|\; Q_n(t_f^{(n)})\right]}{\mathbb{E}\left[T_n[f] \;\middle|\; Q_n(t_f^{(n)})\right]} \tag{11}$$

$$D_n[f] = \frac{V\hat{W}_n(\alpha_n[f])m_{\alpha_n[f]}+Ve_n-Q_n(t_f^{(n)})\mu_n+\mathbb{E}\left[V\hat{g}(\alpha_n[f])I_n[f]+\frac{B_0}{2}T_n[f](T_n[f]-1)\;\middle|\; Q_n(t_f^{(n)})\right]}{\mathbb{E}\left[T_n[f] \;\middle|\; Q_n(t_f^{(n)})\right]} \tag{12}$$

$$D_n[f] = \frac{V\hat{W}_n(\alpha_n[f])m_{\alpha_n[f]}+Ve_n-Q_n(t_f^{(n)})\mu_n+\frac{B_0}{2}\sigma^2_{\alpha_n[f]}+V\hat{g}(\alpha_n[f])I_n[f]}{I_n[f]+m_{\alpha_n[f]}+1}+\frac{B_0}{2}(I_n[f]+m_{\alpha_n[f]}+1). \tag{13}$$

$$\left\lceil\sqrt{\left\lfloor\frac{B_0}{2}\left(Vm_{\alpha_n}\left(\hat{W}_n(\alpha_n)-g_n(\alpha_n)\right)+V\left(e_n-g_n(\alpha_n)\right)-U_n(t_f^{(n)})\mu_n+\frac{B_0}{2}\sigma^2_{\alpha_n}\right)\right\rfloor_{(m_{\alpha_n}+1)^2}}-m_{\alpha_{\alpha_n}}-1\right\rceil^{I_{\max}} \tag{14}$$

$R_n^*(t)$ incoming requests to server n and rejects $d^*(t)$ requests, both of which are random functions on $(\lambda(t), c(t))$. They satisfy the same instantaneous relation as (6). Meanwhile, server n chooses a frame based stationary randomized service decision $(\alpha_n^*[f], I_n^*[f])$, so that the optimal service rate is achieved.

If one knows the stationary distribution for $(\lambda(t), c(t))$, then, this optimal control algorithm can be computed using linear programming. Moreover, the optimal setup cost $W_n^*(t)$, idle cost $g_n^*(t)$, and the active state indicator $H^*(t)$ can also be deduced. Since the algorithm is stationary, these three cost processes are all ergodic Markov processes. Let $T_n^*[f]$ be the frame length process under this algorithm. Thus, it follows from renewal reward theorem that $\left\{\sum_{t=t_f^{(n)}}^{t_{f+1}^{(n)}-1} W_n^*(t)\right\}_{f=0}^{+\infty}$, $\left\{\sum_{t=t_f^{(n)}}^{t_{f+1}^{(n)}-1} g_n^*(t)\right\}_{f=0}^{+\infty}$, $\left\{\sum_{t=t_f^{(n)}}^{t_{f+1}^{(n)}-1} e_nH_n^*(t)\right\}_{f=0}^{+\infty}$, $\left\{\sum_{t=t_f^{(n)}}^{t_{f+1}^{(n)}-1} \mu_n(t)H_n^*(t)\right\}_{f=0}^{+\infty}$ and $\{T_n^*[f]\}_{f=0}^{+\infty}$ are all *i.i.d. random variables over frames*. Let $\overline{C}^*$, $\overline{W}_n^*$, $\overline{G}_n^*$ and $\overline{E}_n^*$ be the optimal time average costs. Let $\overline{R}_n^*$, $\overline{\mu}_n^*$ and $\overline{d}^*$ be the optimal time average routing rate, service rate and rejection rate respectively. Then, by strong law of large numbers,

$$\overline{W}_n^* = \frac{\mathbb{E}\left[\sum_{t=t_f^{(n)}}^{t_f^{(n)}+T_n^*[f]-1} W_n^*(t)\right]}{\mathbb{E}\left[T_n^*[f]\right]} \tag{15}$$

$$\overline{E}_n^* = \frac{\mathbb{E}\left[\sum_{t=t_f^{(n)}}^{t_f^{(n)}+T_n^*[f]-1} e_nH_n^*(t)\right]}{\mathbb{E}\left[T_n^*[f]\right]} \tag{16}$$

$$\overline{G}_n^* = \frac{\mathbb{E}\left[\sum_{t=t_f^{(n)}}^{t_f^{(n)}+T_n^*[f]-1} g_n^*(t)\right]}{\mathbb{E}\left[T_n^*[f]\right]} \tag{17}$$

$$\overline{\mu}_n^* = \frac{\mathbb{E}\left[\sum_{t=t_f^{(n)}}^{t_f^{(n)}+T_n^*[f]-1} \mu_n(t)H_n^*(t)\right]}{\mathbb{E}\left[T_n^*[f]\right]}, \tag{18}$$

Also, notice that $R_n^*(t)$ and $d^*(t)$ depends only on the random variable $\lambda(t)$ and $c(t)$, which is i.i.d. over slots. Thus, $R_n^*(t)$ and $d^*(t)$ are also *i.i.d. random variables over slots*. By law of large numbers,

$$\overline{R}_n^* = \mathbb{E}\left[R_n^*(t)\right], \tag{19}$$

$$\overline{C}^* = \mathbb{E}\left[c(t)d^*(t)\right]. \tag{20}$$

Remark 1: Notice that the idle time $I_n^*[f] \in [1, I_{\max}]$ and the first two moments of the setup time are bounded, it follows the first two moments of $T_n^*[f]$ are bounded.

C. Key features of thresholding algorithm

In this part, we compare the algorithm deduced from the two optimization problems (7) and (11) to that of the best stationary algorithm in section IV-B, illustrating the key features of the proposed online algorithm. Define $\mathcal{F}(t)$ as the system history up till slot t, which includes all the decisions taken and all the random events before slot t. First consider (7). For simplicity of notations, define two random processes $\{X_n[f]\}_{f=0}^{\infty}$ and $\{Z[t]\}_{t=0}^{\infty}$ as follows

$$\begin{aligned} X_n[f] = &\sum_{t=t_f^{(n)}}^{t=t_{f+1}^{(n)}-1}\left(V\left(W_n(t)-\overline{W}_n^*\right)+V\left(e_nH_n(t)-\overline{E}_n^*\right)\right.\\ &+V\left(g_n(t)-\overline{G}_n^*\right)-Q_n(t_f^{(n)})\left(\mu_nH_n(t)-\overline{\mu}^*\right)\\ &\left.+\left(t-t_f^{(n)}\right)B_0-\Psi_n\right),\\ Z[t] = &V\left(c(t)d(t)-\overline{C}^*\right)+\sum_{n=1}^{N}Q_n(t)\left(R_n(t)-\overline{R}_n^*\right), \end{aligned}$$

where $\Psi_n = \frac{B_0}{2}\frac{\mathbb{E}[T_n^*[f](T_n^*[f]-1)]}{\mathbb{E}[T_n^*[f]]}$ and $B_0 = \frac{1}{2}(R_{\max}+\mu_{\max})\mu_{\max}$.

Given the system information $\mathcal{F}(t)$, the random events $c(t)$ and $\lambda(t)$, the solutions (8) and (9) take rejecting and routing decisions so as to minimize (7) over all possible routing and rejecting decisions at time slot t. Thus, the proposed algorithm achieves smaller value on (7) compared to that of the best stationary algorithm in section IV-B. Formally, this idea can be stated as the following inequality: $\mathbb{E}\left[Vc(t)d(t)+\sum_{n=1}^{N}Q_n(t)R_n(t) \;\middle|\; c(t), \lambda(t), \mathcal{F}(t)\right] \leq$

$\mathbb{E}\left[Vc(t)d^*(t) + \sum_{n=1}^{N} Q_n(t)R_n^*(t) \;\middle|\; c(t), \lambda(t), \mathcal{F}(t)\right]$. Taking expectation regarding $c(t)$ and $\lambda(t)$ using the fact that the best stationary algorithm on $R_n^*(t)$ and $d^*(t)$ are i.i.d. over slots (independent of $\mathcal{F}(t)$), together with (19) and (20), we get

$$\mathbb{E}\left[Z(t) \mid \mathcal{F}(t)\right] \leq 0. \tag{21}$$

Similarly, for (11), the proposed service decisions within frame f minimizes $D_n[f]$ in (11), thus, compared to that of the best stationary policy, (22) holds. Again, using the fact that the optimal stationary algorithm gives i.i.d. $W_n^*(t)$, $g_n^*(t)$, $H_n^*(t)$ and $T_n^*[f]$ over frames (independent of $\mathcal{F}(t_f^{(n)})$), as well as (15), (16) and (18), we get

$$\mathbb{E}\left[X_n[f] \;\middle|\; \mathcal{F}(t_f^{(n)})\right] \Big/ \mathbb{E}\left[T_n[f] \;\middle|\; \mathcal{F}(t_f^{(n)})\right] \leq 0 \tag{23}$$

D. Bounded average of supermartingale difference sequeces

The key inequalities (21) and (23) provide us with bounds on the expectations. The following lemma serves as a stepping stone passing from expectation bounds to probability 1 bounds.

Lemma 4: Under the proposed algorithm, the following hold with probability 1,

$$\limsup_{F\to\infty} \frac{1}{F} \sum_{f=0}^{F-1} X_n[f] \leq 0, \tag{24}$$

$$\limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} Z[t] \leq 0. \tag{25}$$

Proof: The key to the proof is treating these two sequences as super-martingale difference sequence and applying law of large numbers for super-martingale difference sequence (theorem 4.1 and corollary 4.2 in [19]).

We first look at the sequence $\{X_n[f]\}_{f=0}^{\infty}$. Let $Y_n[F] = \sum_{f=0}^{F-1} X_n[f]$. We first prove that $Y_n[F]$ is a supermartingale. Notice that $Y_n[F] \in \mathcal{F}\left(t_n^{(F)}\right)$, i.e. it is measurable given all the information before frame $t_n^{(F)}$, and $|Y_n[F]| < \infty$, $\forall F < \infty$. Furthermore, $\mathbb{E}\left[Y_n[F+1] - Y_n[F] \;\middle|\; \mathcal{F}\left(t_F^{(n)}\right)\right] = \mathbb{E}\left[X_n[F] \;\middle|\; \mathcal{F}\left(t_F^{(n)}\right)\right] \leq 0 \cdot \mathbb{E}\left[T_n[F] \;\middle|\; \mathcal{F}\left(t_F^{(n)}\right)\right] = 0$, where the only inequality follows from (23). Thus, it follows $Y_n[F]$ is a supermartingale. Next, we show that the second moment of supermartingale differences, i.e. $\mathbb{E}\left[X_n[f]^2\right]$, is deterministically bounded by a fixed constant for any f. This part of proof is given in Appendix B. Thus, the following holds: $\sum_{f=1}^{\infty} \mathbb{E}\left[X_n[f]^2\right]/f^2 < \infty$. Now, applying corollary 4.2 in [19] immediately gives (24).

Similarly, we can prove (25) by proving $M[t] = \sum_{t=0}^{T-1} Z[t]$ is a supermartingale with bounded second moment on differences using (19), (20) and (21). The procedure is almost the same as above and we omitted the details here for brevity. ■

Corollary 1: The following ratio of time averages is upper bounded with probability 1,
$\limsup_{F\to\infty} \sum_{f=0}^{F-1} X_n[f] \Big/ \sum_{f=0}^{F-1} T_n[f] \leq 0$.

Proof: From (24), it follows for any $\epsilon > 0$, there exists an $F_0(\epsilon)$ such that $F \geq F_0(\epsilon)$ implies $\sum_{f=0}^{F-1} X_n[f] \Big/ \sum_{f=0}^{F-1} T_n[f] \leq \epsilon \Big/ \frac{1}{F}\sum_{f=0}^{F-1} T_n[f] \leq \epsilon$. Thus, $\limsup_{F\to\infty} \sum_{f=0}^{F-1} X_n[f] \Big/ \sum_{f=0}^{F-1} T_n[f] \leq \epsilon$. Since ϵ is arbitrary, take $\epsilon \to 0$ gives the result. ■

E. Near optimal time average cost

The ratio of time averages in corollary 1 and the true time average share the same bound, which is proved by the following lemma:

Lemma 5: The following time average is bounded,

$$\begin{aligned}
&\limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \Big(V\left(W_n(t) + e_n H_n(t) + g_n(t)\right) - Q_n(t_f^{(n)}) \\
&(\mu_n H_n(t) - \overline{\mu}_n^*) + \left(t - t_f^{(n)}\right) B_0\Big) \\
&\leq V(\overline{W}_n^* + \overline{E}_n^* + \overline{G}_n^*) + \Psi_n,
\end{aligned} \tag{26}$$

where $\Psi_n = \frac{B_0}{2} \frac{\mathbb{E}[T_n^*[f](T_n^*[f]-1)]}{\mathbb{E}[T_n^*[f]]}$ and $B_0 = \frac{1}{2}(R_{\max} + \mu_{\max})\mu_{\max}$.

The idea of the proof is similar to that of basic renewal theory, which derives upper and lower bounds for each T within any frame F using corollary 1, thereby showing that as $T \to \infty$, the upper and lower bounds meet. See appendix C for details. With this lemma, we have:

Theorem 1: If $Q_n(0) = 0, \forall n \in \mathcal{N}$, then the time average total cost under the algorithm is near optimal on the order of $\mathcal{O}(1/V)$, i.e.

$$\begin{aligned}
&\limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \left(c(t)d(t) + \sum_{n=1}^{N} \left(W_n(t) + e_n H_n(t) + g_n(t)\right)\right) \\
&\leq \underbrace{\overline{C}^* + \sum_{n=1}^{N} \left(\overline{W}_n^* + \overline{E}_n^* + \overline{G}_n^*\right)}_{\text{optimal cost}} + \frac{\sum_{n=1}^{N} \Psi_n + B_3}{V},
\end{aligned} \tag{27}$$

where $B_3 \triangleq \frac{1}{2}\sum_{n=1}^{N}(R_{\max}+\mu_n)^2$, $\Psi_n = \frac{B_0}{2} \frac{\mathbb{E}[T_n^*[f](T_n^*[f]-1)]}{\mathbb{E}[T_n^*[f]]}$ and $B_0 = \frac{1}{2}(R_{\max} + \mu_{\max})\mu_{\max}$.
The proof uses Lyapunov drift analysis. See appendix D for details. As parameter V can be chosen arbitrarily large, the final term in (27) can be made arbitrarily small, with a tradeoff on queue size which is linear on V (as per Lemma 2).

V. Delay improvement via virtualization

A. Delay improvement

The algorithm in previous sections optimizes time average cost. However, it can route requests to idle queues, which increases system delay. This section considers an improvement that maintains the same average cost guarantees, but reduces delay. This is done by a "virtualization" technique that reduces from N server request queues to only one request queue $Q(t)$. Specifically, the same Algorithm 1 is run, with queue updates (1) for each of the N queues $Q_n(t)$. However, the $Q_n(t)$ processes are now virtual queues rather than actual queues: Their values are only kept in software. Every slot t, the data center observes the incoming requests $\lambda(t)$, rejection cost $c(t)$ and virtual queue values, making rejection decision according to (8) as before. The admitted requests are queued in $Q(t)$.

$$\frac{\mathbb{E}\left[\sum_{t=t_f^{(n)}}^{t=t_{f+1}^{(n)}-1}\left(V(W_n(t)+e_nH_n(t)+g_n(t))-Q_n(t_f^{(n)})\mu_n(t)H_n(t)\right)+\left(t-t_f^{(n)}\right)B_0 \;\middle|\; \mathcal{F}(t_f^{(n)})\right]}{\mathbb{E}\left[T_n[f] \;\middle|\; \mathcal{F}(t_f^{(n)})\right]}$$
$$\leq \frac{\mathbb{E}\left[\sum_{t=t_f^{(n)}}^{t=t_f^{(n)}+T_n^*[f]-1}\left(V\left(W_n^*(t)+e_nH_n^*(t)+g_n^*(t)\right)-Q_n(t_f^{(n)})\mu_nH_n^*(t)\right)+\frac{B_0}{2}T_n^*[f](T_n^*[f]-1) \;\middle|\; \mathcal{F}(t_f^{(n)})\right]}{\mathbb{E}\left[T_n^*[f] \;\middle|\; \mathcal{F}(t_f^{(n)})\right]} \quad (22)$$

Meanwhile, each server n makes active/idle decisions observing its own virtual queue $Q_n(t)$ same as before. Whenever a server is active, it grabs the requests from request queue $Q(t)$ and serves them. This results in an actual queue updating for the system:

$$Q(t+1)=\max\left\{Q(t)+\lambda(t)-d(t)-\sum_{n=1}^{N}\mu_n(t)H_n(t),\ 0\right\}. \quad (28)$$

Fig. 2 shows this data center architecture.

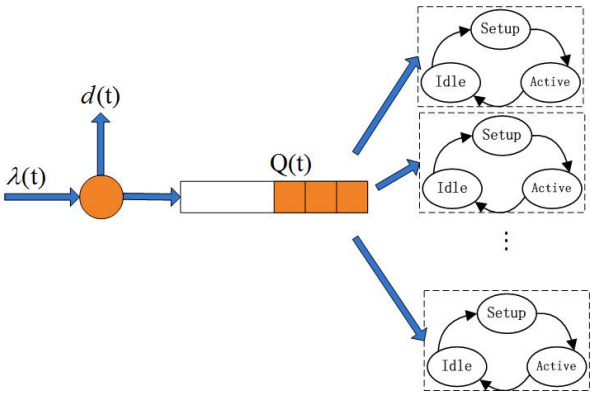

Fig. 2. Illustration of basic data center architecture.

B. Performance guarantee

Since this algorithm does not look at the actual queue $Q(t)$, it is not clear whether or not the actual request queue would be stabilized under the proposed algorithm. The following lemma answers the question. For simplicity, we call the system with N queues, where our algorithm applies, the virtual system, and call the system with only one queue the actual system.

Lemma 6: If $Q(0)=0$ and $Q_n(0)=0,\ \forall n\in\mathcal{N}$, then the virtualization technique stabilizes the queue $Q(t)$ with the bound: $Q(t)\leq N(Vc_{\max}+R_{\max})$.

Proof: Notice that this bound is N times the individual queue bound in lemma 2, we prove the lemma by showing that the sum-up weights $\sum_{n=1}^{N}Q_n(t)$ in the virtual system always dominates the queue length $Q(t)$. We prove this by induction. The base case is obvious since $Q(0)=\sum_{n=1}^{N}Q_n(0)=0$. Suppose at the beginning of time t, $Q(t)\leq\sum_{n=1}^{N}Q_n(t)$, then, during time t, we distinguish between the following two cases:

1) Not all active servers in actual system have requests to serve. This case happens if and only if there are not enough requests in $Q(t)$ to be served, i.e. $\lambda(t)-d(t)+Q(t)<\sum_{n=1}^{N}\mu_n(t)H_n(t)$. Thus, according to queue updating rule (28), at the beginning of time slot $t+1$, there will be no request sitting in the actual queue, i.e. $Q(t+1)=0$. Hence, it is guaranteed that $Q(t+1)\leq\sum_{n=1}^{N}Q_n(t+1)$.
2) All active servers in the actual system have requests to serve. Notice that the virtual system and the actual system have exactly the same arrivals, rejections and server active/idle states. Thus, the following holds, $Q(t+1)=Q(t)+\lambda(t)-d(t)-\sum_{n=1}^{N}\mu_n(t)H_n(t)\leq\sum_{n=1}^{N}Q_n(t)+\sum_{n=1}^{N}R_n(t)-\sum_{n=1}^{N}\mu_n(t)H_n(t)\leq\sum_{n=1}^{N}\max\{Q_n(t)+R_n(t)-\mu_n(t)H_n(t),\ 0\}=\sum_{n=1}^{N}Q_n(t+1)$, where the first inequality follows from induction hypothesis as well as the fact that $\sum_{n=1}^{N}R_n(t)=\lambda(t)-d(t)$.

Overall, we proved $Q(t)\leq\sum_{n=1}^{N}Q_n(t),\ \forall t$. Since each $Q_n(t)\leq Vc_{\max}+R_{\max},\ \forall t$, the lemma follows. ■

Since the virtual system and the actual system have exactly the same cost, and it can be shown that the optimal cost in one queue system is lower bounded by the optimal cost in N queue system, the near optimal performance is still guaranteed.

C. Further improvement by over-provisioning

The workload in data center is highly unpredictable and have short term fluctuations and spikes ([12]). In order to further reduce the delay of requests, the data center can keep some of the servers always on, so that in case of peak traffic, not all servers have to go through the setup state in order to meet the need. Moreover, the data center can do "over-provisioning". Specifically, every time slot t, we add a "virtual request flow" $\lambda_v(t)$ on top of actual requests $\lambda(t)$ based on the following three rules: 1)If rejection decisions occur, virtual requests are rejected instead of actual requests whenever possible. 2)Rejected virtual requests do not generate actual cost, but their virtual costs are still considered for algorithmic decisions, 3) Virtual requests are treated as a low priority requests when being served, i.e. servers will not serve virtual requests until all real requests are served.

Then, we keep Algorithm 1 unchanged except that the reject decision (8) and routing decision (9) will work on virtual requests also. Then, according to the virtual queue update (1), the virtual queue length tends to be larger, which in turn pushes the servers to be active more often. Meanwhile, the algorithm achieves near optimality with respect to the original request rate $\overline{\lambda}$ plus the virtual rate $\overline{\lambda}_v$.

TABLE II
PROBLEM PARAMETERS

Server	μ_n	e_n	$\hat{W}_n(\alpha_n)$	$\mathbb{E}[\hat{\tau}(\alpha_n)]$
1	$\{2,3,4,5,6\}$	4	2	5.893
2	$\{2,3,4\}$	2	3	4.342
3	$\{2,3,4\}$	3	3	27.397
4	$\{1,2,3\}$	4	2	5.817
5	$\{2,3,4\}$	2	4	6.211

VI. SIMULATION

In this section, we demonstrate the near optimality of the proposed algorithm for an N queues system and compare the performances of the proposed algorithm under N queues setting and one queue setting respectively.

A. Near optimality in N queues system

In the first experiment, we set number of servers $N = 5$. The incoming requests $\lambda(t)$ are integers following a uniform distribution in $[10, 30]$. The request rejection costs $c(t)$ are also integers following a uniform distribution in $[1, 6]$. The maximum admission amount $R_{\max} = 40$ and the maximum idle time $I_{\max} = 1000$. There is only one idle option α_n for each server where the idle cost $\hat{g}(\alpha_n) = 0$. The setup time follows a geometric distribution with mean $\mathbb{E}[\hat{\tau}(\alpha_n)]$, setup cost $\hat{W}_n(\alpha_n)$ per slot, service cost e_n per slot, and the service amount μ_n follows a uniform distribution over integers. The values $1/\mathbb{E}[\hat{\tau}(\alpha_n)]$ are generated uniform at random within $[0, 1]$ and specified in table II.

The algorithm is run for 1 million slots in each trial and each plot takes the average of these 1 million slots. We compare our algorithm to the optimal stationary algorithm. The optimal stationary algorithm is computed using linear program [21] with the full knowledge of the statistics of requests and rejecting costs.

In Fig. 3, we show that as our tradeoff parameter V gets larger, the average cost approaches the optimal value and achieves a near optimal performance. Fig. 4 plots the average sum-up queue size $\sum_{n=1}^{5} Q_n(t)$ and shows as V gets larger, the average sum-up queue size becomes larger. We also plot the sum of individual queue bound from Lemma 2 for comparison. The bound is quite loose, indicating that we are getting a much better delay than expected from theorems.

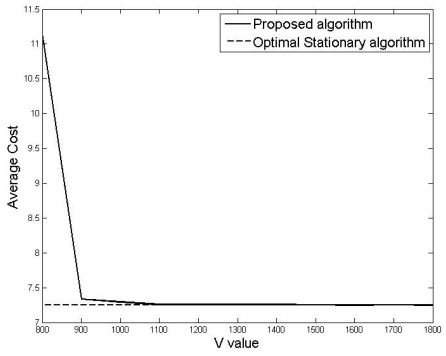

Fig. 3. Time average cost verses V parameter over 1 millon slots.

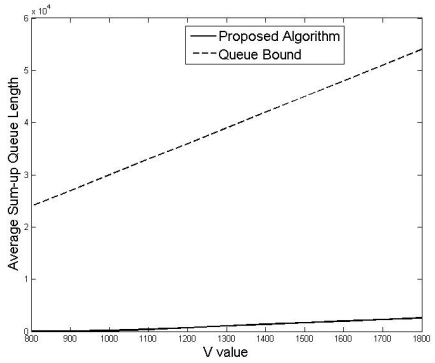

Fig. 4. Time average sum-up request queue length verses V parameter over 1 millon slots.

We then tune the requests $\lambda(t)$ to be uniform in $[20, 40]$ and keep other parameters unchanged. In Fig. 5, we see that since the request rate gets larger, we need V to be larger in order to obtain the near optimality. Fig. 6 gives the sum-up average queue length under this case. The average queue length is larger than that of Fig. 4.

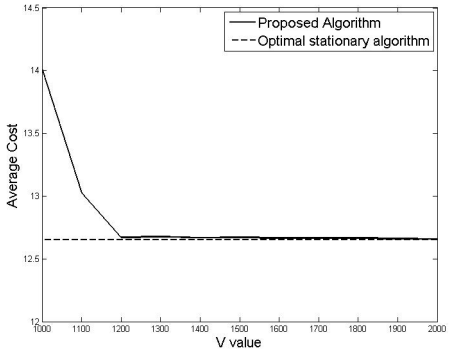

Fig. 5. Time average cost verses V parameter over 1 millon slots.

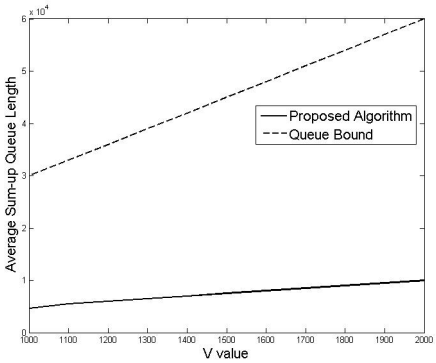

Fig. 6. Time average cost verses V parameter over 1 millon slots.

B. Data center with unpredictable peak traffic

This section considers the data center with single queue improvement from section V. Specifically, we consider a data

center with 20 homogeneous servers with integer service rate $\mu(t)$ following a uniform distribution between $[2, 8]$, service cost $e = 50$ per slot. There is one idle option with idle cost $g = 0$, setup time geometrically distributed with success probability $\rho = 0.01$(take 100 slots to setup on average) and setup cost $W = 50$ per slot. The rejecting cost $c(t)$ is uniform between $[10, 20]$ per request. The maximum admission amount $R_{\max} = 200$ and the maximum idle time $I_{\max} = 1000$. Meanwhile, we fix the trade-off parameter $V = 7000$ and keep 5 servers always active.

We simulate above system over 150 thousand slots. In order to capture the unpredictability and large fluctuations of data center traffic flows, the incoming requests are designed as follows: The "base" incoming request traffic is uniform between $[10, 30]$ per slot. We then place 5 peaks on random places within these 150 thousand slots. Each peak lasts for 5000 slots and has the peak incoming requests uniform between $[50, 150]$ per slot.

Fig. 7 shows the sample path of incoming traffic over 150 thousand slots. Fig. 8 plots the corresponding running average as well as parameter $\alpha = 0.001$ exponential average [1] of overall costs in data center. The exponential average emphasizes the local performance of the algorithm whereas the running average gives the overall performance. For comparison, we also plotted the optimal cost with incoming rate $\lambda = 20$ per slot and $\lambda = 100$ per slot respectively. It can be seen from the exponential average plot that the local average cost results from the proposed algorithm tracks the change of incoming request traffic quite closely.

Fig. 9 plots the sample path of actual queue length $Q(t)$ and sum-up virtual queue length $\sum_{n=1}^{20} Q_n(t)$. It can be seen that the sum-up virtual queue length always dominates the actual queue length, which coincides with lemma 6. Furthermore, Fig. 9 also demonstrates experimentally that building up request queues on servers' side might result in much worse delay than keeping only one system request queue. Notice that the peak queue length is around 50000 requests whereas the processing rate is 100 requests per slot all out. Thus, the peak delay would be around 500 slots, which is roughly in the order of expected setup time (100 slots).

Finally, Fig. 10 plots the number of active servers verses slot number under the proposed algorithm. It can be seen that the algorithm quickly adapts to the new workload by increasing or decreasing the number of active servers even when the workload changes in an unpredictable manner.

C. Over-provisioning and delay improvement

This section further analyzes the effects of over provision on request delay. We keep all the parameters the same as previous experiment except that we do over-provisioning by adding a virtual request flow. Fig. 11 and Fig. 12 plot the exponential time average costs and sample paths of $Q(t)$ under over provision with $\lambda_v(t) = \lambda(t)$ and no over provision.

We can see from Fig. 12 that over-provisioning drastically reduces delay. This is because virtual requests make virtual queues larger, which inclines servers to stay active after a peak wave of traffic. Thus, servers do not have to go through a setup time when another peak occurs shortly after. However, from Fig. 11 we see that such over-provisioning increases cost.

[1]For a sequence of real values $\{X_t\}_{t=0}^{\infty}$, the exponential average at time slot T (denoted as $\widetilde{X}_T$) is $\widetilde{X}_T = (1-\alpha)\widetilde{X}_{T-1} + \alpha X_T$.

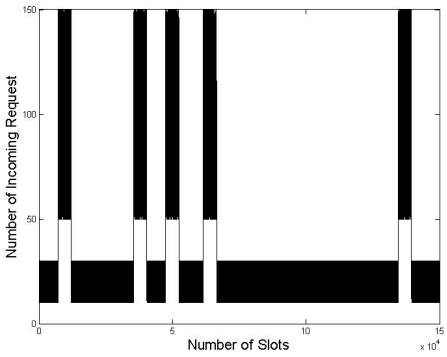

Fig. 7. Sample path of incoming request rate.

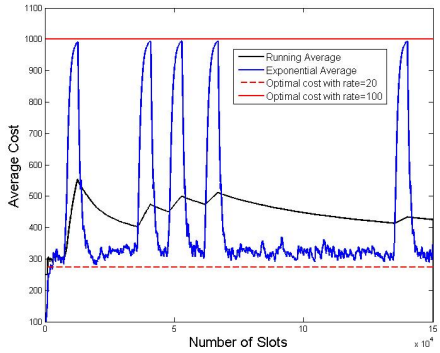

Fig. 8. Running average and exponential average costs over 150 thousand slots.

APPENDIX A— PROOF OF LEMMA 1

We first show that the optimal decision on $I_n[f]$ is a pure decision. We have (29) holds, where the first equality follows from the definition $T_n[f] = I_n[f] + \tau_n[f] + 1$ and the second equality follows from the fact that the setup time

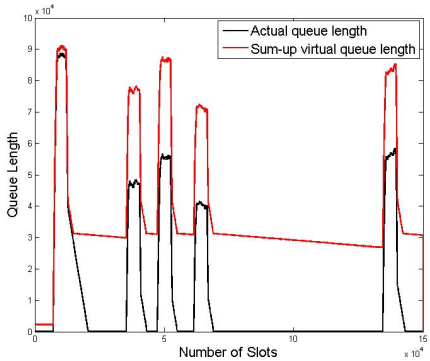

Fig. 9. Sample pathes of actual queue length and sum-up virtual queue length.

$$D_n[f] = \frac{V\hat{W}_n(\alpha_n[f])m_{\alpha_n[f]} + Ve_n - Q_n(t_f^{(n)})\mu_n + \mathbb{E}\left[\frac{B_0}{2}(I_n[f]+\tau_n[f]+1)^2 + V\hat{g}(\alpha_n[f])I_n[f] \,\middle|\, Q_n(t_f^{(n)})\right]}{\mathbb{E}\left[I_n[f]+\tau_n[f]+1 \,\middle|\, Q_n(t_f^{(n)})\right]} - \frac{B_0}{2}$$
$$= \frac{V\hat{W}_n(\alpha_n[f])m_{\alpha_n[f]} + Ve_n - Q_n(t_f^{(n)})\mu_n + \mathbb{E}\left[\frac{B_0}{2}(I_n[f]+m_{\alpha_n}+1)^2 + \frac{B_0}{2}\sigma^2_{\alpha_n[f]} + V\hat{g}(\alpha_n[f])I_n[f] \,\middle|\, Q_n(t_f^{(n)})\right]}{\mathbb{E}\left[I_n[f]+m_{\alpha_n}+1 \,\middle|\, Q_n(t_f^{(n)})\right]} - \frac{B_0}{2} \tag{29}$$

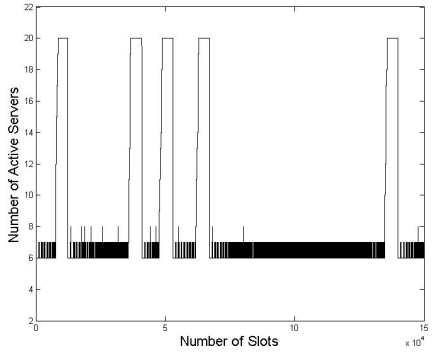

Fig. 10. Number of active servers over time.

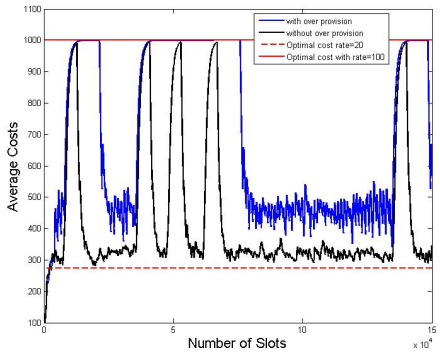

Fig. 11. Exponential average costs with and without over provision.

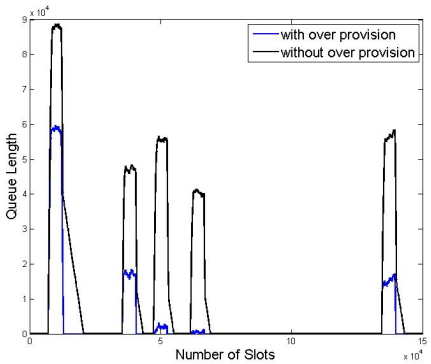

Fig. 12. Sample paths of queue lengths with and without over provision.

$\tau_n[f] = \hat{\tau}_n(\alpha_n[f])$ is independent of $I_n[f]$ with mean $m_{\alpha_n[f]}$ and variance $\sigma^2_{\alpha_n[f]}$. For simplicity of notations, let

$$\begin{aligned} F(\alpha_n[f], I_n[f]) =& V\hat{W}_n(\alpha_n[f])m_{\alpha_n[f]} + Ve_n - Q_n(t_f^{(n)})\mu_n \\ &+ \frac{B_0}{2}(I_n[f]+m_{\alpha_n[f]}+1)^2 + V\hat{g}(\alpha_n[f])I_n[f] \\ &+ \frac{B_0}{2}\sigma^2_{\alpha_n[f]} \\ G(\alpha_n[f], I_n[f]) =& I_n[f] + m_{\alpha_n[f]} + 1, \end{aligned}$$

then

$$D_n[f] = \frac{\mathbb{E}\left[F(\alpha_n[f], I_n[f]) \mid Q_n(t_f^{(n)})\right]}{\mathbb{E}\left[G(\alpha_n[f], I_n[f]) \mid Q_n(t_f^{(n)})\right]} - \frac{B_0}{2}.$$

Meanwhile, given the queue length $Q_n(t_f^{(n)})$ at frame f, denote the benchmark solution over pure decisions as

$$m \triangleq \min_{I_n[f]\in\mathbb{N},\ I_n[f]\in[1,I_{\max}],\alpha_n[f]\in\mathcal{L}_n} \frac{F(\alpha_n[f], I_n[f])}{G(\alpha_n[f], I_n[f])}. \tag{30}$$

Then, for any randomized decision on $\alpha_n[f]$ and $I_n[f]$, its realization within frame f satisfies the following

$$\frac{F(\alpha_n[f], I_n[f])}{G(\alpha_n[f], I_n[f])} \geq m,$$

which implies

$$F(\alpha_n[f], I_n[f]) \geq mG(\alpha_n[f], I_n[f]).$$

Taking conditional expectation from both sides gives

$$\begin{aligned} &\mathbb{E}\left[F(\alpha_n[f], I_n[f]) \mid Q_n(t_f^{(n)})\right] \\ &\geq m\mathbb{E}\left[G(\alpha_n[f], I_n[f]) \mid Q_n(t_f^{(n)})\right] \\ &\Rightarrow \frac{\mathbb{E}\left[F(\alpha_n[f], I_n[f]) \mid Q_n(t_f^{(n)})\right]}{\mathbb{E}\left[G(\alpha_n[f], I_n[f]) \mid Q_n(t_f^{(n)})\right]} \geq m. \end{aligned}$$

Thus, it is enough to consider pure decisions only, which boils down to computing (30). This proves the lemma.

Appendix B— Proof of Lemma 4

This section is dedicated to prove that $\mathbb{E}\left[X_n[f]^2\right]$ is bounded. First of all, since the idle option set $\mathcal{L}_n$ is finite, denote

$$W_{\max} = \max_{\alpha_n\in\mathcal{L}_n} W_n(\alpha_n)$$
$$g_{\max} = \max_{\alpha_n\in\mathcal{L}_n} g_n(\alpha_n)$$

It is obvious that $|W_n(t)-\overline{W}_n^*| \le W_{\max}$, $|g_n(t)-\overline{G}_n^*| \le g_{\max}$, $|e_nH_n(t)-\overline{E}_n^*| \le e_n$, and $|\mu_nH_n(t)-\overline{\mu}^*| \le \mu_n$. Combining with the boundedness of queues in lemma 2, it follows

$$\begin{aligned}|X_n[f]| \le & \sum_{t=t_f^{(n)}}^{t=t_{f+1}^{(n)}-1} \left(V\left(W_{\max}+e_n+g_{\max}\right)+\left(Vc_{\max}\right.\right.\\ &\left.\left.+R_{\max}\right)\mu_n+\left(t-t_f^{(n)}\right)B_0+\Psi_n\right)\\ \le & \left(V(W_{\max}+e_n+g_{\max})+(Vc_{\max}+R_{\max})\mu_n\right.\\ &\left.+\Psi_n\right)T_n[f]+\frac{T_n[f](T_n[f]-1)B_0}{2}\end{aligned}$$

Let $B_1 \triangleq V(W_{\max}+e_n+g_{\max})+(Vc_{\max}+R_{\max})\mu_n+\Psi_n+B_0/2$, it follows

$$|X_n[f]| \le B_1T_n[f]+\frac{B_0}{2}T_n[f]^2.$$

Thus,

$$\mathbb{E}\left[X_n[f]^2\right] \le B_1^2\mathbb{E}\left[T_n[f]^2\right]+B_1B_0\mathbb{E}\left[T_n[f]^3\right]+\frac{B_0^2}{4}\mathbb{E}\left[T_n[f]^4\right].$$

Notice that $T_n[f] \le I_n[f]+\tau_n[f]+1$ by (2), where $I_n[f]$ is upper bonded by $I_{\max}$ and $\tau_n[f]$ has first four moments bounded by assumption 2. Thus, $\mathbb{E}\left[X_n[f]^2\right]$ is bounded by a fixed constant.

Appendix C— Proof of Lemma 5

Proof: Let's first abbreviate the notation by defining

$$\begin{aligned}Y(t) =&V(W_n(t)+e_nH_n(t)+g_n(t))-Q_n(t_f^{(n)})(\mu_nH_n(t)-\overline{\mu}_n^*)\\ &+\left(t-t_f^{(n)}\right)B_0.\end{aligned}$$

For any $T \in [t_F^{(n)},\ t_F^{(n+1)})$, we can bound the partial sums from above by the following

$$\sum_{t=0}^{T-1}Y(t) \le \sum_{t=0}^{t_F^{(n)}-1}Y(t)+B_2T_n[F]+\frac{B_0}{2}T_n[F]^2,$$

where $B_0=\frac{1}{2}(R_{\max}+\mu_{\max})\mu_{\max}$ is defined in (11), and $B_2 \triangleq VW_n+V\mu_ne_n+(Vc_{\max}+R_{\max})\mu_n+B_0/2$. Thus,

$$\begin{aligned}\frac{1}{T}\sum_{t=0}^{T-1}Y(t) &\le \frac{1}{T}\sum_{t=0}^{t_F^{(n)}-1}Y(t)+\frac{1}{T}\left(B_2T_n[F]+\frac{B_0}{2}T_n[F]^2\right)\\ &\le \max\{a[F],\ b[F]\},\end{aligned}$$

where

$$a[F] \triangleq \frac{1}{t_F^{(n)}}\sum_{t=0}^{t_F^{(n)}-1}Y(t)+\frac{1}{t_F^{(n)}}\left(B_2T_n[F]+\frac{B_0}{2}T_n[F]^2\right),$$

$$b[F] \triangleq \frac{1}{t_{F+1}^{(n)}}\sum_{t=0}^{t_F^{(n)}-1}Y(t)+\frac{1}{t_{F+1}^{(n)}}\left(B_2T_n[F]+\frac{B_0}{2}T_n[F]^2\right).$$

Thus, this implies that

$$\begin{aligned}\limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}Y(t) &\le \limsup_{F\to\infty}\max\{a[F],\ b[F]\}\\ &= \max\left\{\limsup_{F\to\infty}a[F],\ \limsup_{F\to\infty}b[F]\right\}.\end{aligned}$$

We then try to work out an upper bound for $\limsup_{F\to\infty}a[F]$ and $\limsup_{F\to\infty}b[F]$ respectively.

1) Bound for $\limsup_{F\to\infty}a[F]$:

$$\begin{aligned}\limsup_{F\to\infty}a[F] \le &\limsup_{F\to\infty}\frac{1}{t_F^{(n)}}\sum_{t=0}^{t_F^{(n)}-1}Y(t)\\ &+\limsup_{F\to\infty}\frac{1}{t_F^{(n)}}\left(B_2T_n[F]+\frac{B_0}{2}T_n[F]^2\right)\\ \le &V(\overline{W}_n^*+\overline{E}_n^*+\overline{G}_n^*)+\Psi_n\\ &+\limsup_{F\to\infty}\frac{1}{t_F^{(n)}}\left(B_2T_n[F]+\frac{B_0}{2}T_n[F]^2\right).\end{aligned}$$

where the second inequality follows from corollary 1. It remains to show that

$$\limsup_{F\to\infty}\frac{1}{t_F^{(n)}}\left(B_2T_n[F]+\frac{B_0}{2}T_n[F]^2\right) \le 0. \quad (31)$$

Since $t_F^{(n)} \ge F$, it is enough to show that

$$\limsup_{F\to\infty}\frac{T_n[F]}{F}=0, \quad (32)$$

$$\limsup_{F\to\infty}\frac{T_n[F]^2}{F}=0. \quad (33)$$

We prove (33), and (32) is similar. Since each $T_n[F]=I_n[F]+\tau_n[F]+1$, where $I_n[F] \le I_{\max}$ and $\tau_n[F]$ has bounded first four moments, the first four moments of $T_n[F]$ must also be bounded and there exists a constant $C>0$ such that

$$\mathbb{E}\left[T_n[F]^4\right] \le C.$$

For any $\epsilon>0$, define a sequence of events

$$A_F^\epsilon \triangleq \left\{T_n[F]^2 > \epsilon F\right\}.$$

According to Markov inequality,

$$Pr\left[A_F^\epsilon\right] \le \frac{\mathbb{E}\left[T_n[F]^4\right]}{\epsilon^2F^2} \le \frac{C}{\epsilon^2F^2}.$$

Thus,

$$\sum_{F=1}^{\infty}Pr\left[A_F^\epsilon\right] \le \frac{C}{\epsilon^2}\sum_{F=1}^{\infty}\frac{1}{F^2} \le \frac{2C}{\epsilon^2} < \infty.$$

By Borel-Cantelli lemma (lemma 1.6.1 in [20]),

$$Pr\left[A_F^\epsilon \text{ occurs infinitely often}\right]=0,$$

which implies

$$Pr\left[\limsup_{F\to\infty}\frac{T_n[F]^2}{F} > \epsilon\right]=0.$$

Since ϵ is arbitrary, this implies (33). Similarly, (32) can be proved. Thus, (31) holds and

$$\limsup_{F\to\infty}a[F] \le V(\overline{W}_n^*+\overline{E}_n^*+\overline{G}_n^*)+\Psi_n.$$

2) Bound for $\limsup_{F\to\infty} b[F]$:

$$\begin{aligned}\limsup_{F\to\infty} b[F] \leq & \limsup_{F\to\infty} \frac{1}{t_F^{(n)}} \sum_{t=0}^{t_F^{(n)}-1} Y(t) \cdot \frac{t_F^{(n)}}{t_{F+1}^{(n)}} \\ & + \limsup_{F\to\infty} \frac{1}{t_{F+1}^{(n)}} \left(B_2 T_n[F] + \frac{B_0}{2} T_n[F]^2 \right). \\ \leq & \limsup_{F\to\infty} \left(\frac{1}{t_F^{(n)}} \sum_{t=0}^{t_F^{(n)}-1} Y(t) \right) \\ & \cdot \limsup_{F\to\infty} \frac{t_F^{(n)}}{t_{F+1}^{(n)}} \\ \leq & \left(V\left(\overline{W}_n^* + \overline{E}_n^* + \overline{G}_n^* \right) + \Psi_n \right) \\ & \cdot \limsup_{F\to\infty} \frac{t_F^{(n)}}{t_{F+1}^{(n)}} \\ \leq & V\left(\overline{W}_n^* + \overline{E}_n^* + \overline{G}_n^* \right) + \Psi_n,\end{aligned}$$

where the second inequality follows from (31), the third inequality follows from corollary 1 and the last inequality follows from the fact that $V\left(\overline{W}_n^* + \overline{E}_n^* + \overline{G}_n^* \right) + \Psi_n > 0$.

Overall, we proved the lemma. ■

Appendix D— Proof of Theorem 1

Proof: Define the drift-plus-penalty(DPP) expression $P(t)$ as follows

$$\begin{aligned}P(t) = & V\left(c(t)d(t) + \sum_{n=1}^{N} \left(W_n(t) + e_n H_n(t) + g_n(t) \right) \right) \\ & + \frac{1}{2} \sum_{n=1}^{N} \left(Q_n(t+1)^2 - Q_n(t)^2 \right).\end{aligned}$$

By simple algebra using the queue updating rule (1), we can work out the upper bound for $P(t)$ as follows,

$$\begin{aligned}P(t) \leq & \frac{1}{2} \sum_{n=1}^{N} (R_n(t) + \mu_n)^2 + V\left(c(t)d(t) + \right. \\ & \left. \sum_{n=1}^{N} \left(W_n(t) + e_n H_n(t) + g_n(t) \right) \right) \\ & + \sum_{n=1}^{N} Q_n(t)(R_n(t) - \mu_n H_n(t)) \\ \leq & B_3 + V\left(c(t)d(t) + \sum_{n=1}^{N} \left(W_n(t) + e_n H_n(t) + g_n(t) \right) \right) \\ & + \sum_{n=1}^{N} Q_n(t)(R_n(t) - \mu_n H_n(t)) \\ \leq & B_3 + V c(t) d(t) + \sum_{n=1}^{N} Q_n(t) \left(R_n(t) - \overline{R}_n^* \right) \\ & + V \sum_{n=1}^{N} \left(W_n(t) + e_n H_n(t) + g_n(t) \right) \\ & + \sum_{n=1}^{N} Q_n(t) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right)\end{aligned}$$

where $B_3 = \frac{1}{2} \sum_{n=1}^{N} (R_{\max} + \mu_n)^2$, the last inequality follows from adding $\sum_{n=1}^{N} Q_n(t) \overline{\mu}_n^*$ and subtracting $\sum_{n=1}^{N} Q_n(t) \overline{R}^*$ with the fact that the best randomized stationary algorithm should also satisfy the constraint (5), i.e. $\overline{\mu}^* \geq \overline{R}_n^*$.

Now we take the partial average of $P(t)$ from 0 to $T-1$ and take $\limsup_{T\to\infty}$,

$$\begin{aligned}& \limsup_{T\to\infty} \frac{1}{T} \sum_{n=1}^{T-1} P(t) \\ \leq & B_3 + \limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \left(V c(t) d(t) + \sum_{n=1}^{N} Q_n(t) \left(R_n(t) - \overline{R}_n^* \right) \right) \\ & + \sum_{n=1}^{N} \limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \left(V \left(W_n(t) + e_n H_n(t) + g_n(t) \right) \right. \\ & \left. + Q_n(t) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) \right). \end{aligned} \tag{34}$$

According to (25),

$$\limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \left(V c(t) d(t) + \sum_{n=1}^{N} Q_n(t) \left(R_n(t) - \overline{R}_n^* \right) \right) \leq V \overline{C}^*. \tag{35}$$

On the other hand,

$$\begin{aligned}& \limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \left(V \left(W_n(t) + e_n H_n(t) + g_n(t) \right) \right. \\ & \left. + Q_n(t) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) \right) \\ \leq & \limsup_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T-1} \left(V \left(W_n(t) + e_n H_n(t) + g_n(t) \right) \right. \\ & \left. + Q_n(t_f^{(n)}) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) + (t - t_f^{(n)}) B_0 \right) \\ \leq & V\left(\overline{W}_n^* + \overline{E}_n^* + \overline{G}_n^* \right) + \Psi_n, \end{aligned} \tag{36}$$

where $B_0 = \frac{1}{2}(R_{\max} + \mu_{\max})\mu_{\max}$ as defined below (11), the first inequality follows from the fact that for any $t \in \left(t_f^{(n)},\ t_{f+1}^{(n)} \right)$,

$$\begin{aligned}& Q_n(t) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) \\ \leq & Q_n(t_f^{(n)}) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) \\ & + \left(Q_n(t) - Q_n(t_f^{(n)}) \right) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) \\ \leq & Q_n(t_f^{(n)}) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) \\ & + \sum_{t=t_f^{(n)}}^{t_{f+1}^{(n)}-1} \left(R_n(t) - \mu_n H_n(t) \right) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) \\ \leq & Q_n(t_f^{(n)}) \left(\overline{\mu}_n^* - \mu_n H_n(t) \right) + (t - t_f^{(n)}) B_0,\end{aligned}$$

and the second inequality follows from lemma 5. Substitute (35) and (36) into (34) gives

$$\limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}P(t)\leq V\left(\overline{C}^*+\sum_{n=1}^{N}\left(\overline{W}_n^*+\overline{E}_n^*+\overline{G}_n^*\right)\right)+B_3+\sum_{n=1}^{N}\Psi_n. \quad (37)$$

Finally, notice that by telescoping sums,

$$\begin{aligned}&\limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}P(t)\\ =&\limsup_{T\to\infty}\left(\frac{V}{T}\sum_{t=0}^{T-1}\left(c(t)d(t)+\sum_{n=1}^{N}\left(W_n(t)+e_nH_n(t)\right.\right.\right.\\ &\left.\left.\left.+g_n(t)\right)\right)+\frac{1}{2}\sum_{n=1}^{N}Q_n(T)^2\right)\\ \geq& V\cdot\limsup_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\left(c(t)d(t)+\sum_{n=1}^{N}\left(W_n(t)+e_nH_n(t)\right.\right.\\ &\left.\left.+g_n(t)\right)\right)\end{aligned}$$

Substituting the above inequality into (37) and divide V from both sides give the desired result. ■

REFERENCES

[1] D. P. Bertsekas. *Dynamic Programming and Optimal Control, vols. 1 and 2*. Athena Scientific, Belmont, Mass, 1995.

[2] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* John Wiley & Sons, 2005.

[3] S. Ross. *Introduction to Probability Models.* Academic Press, 8th edition, 2002.

[4] R. Gallager. *Discrete Stochastic Processes. Kluwer Academic Publishers*, Boston, 1996.

[5] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1-149, 2006.

[6] M. J. Neely, Dynamic Optimization and Learning for Renewal Systems, *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 32-46, Jan. 2013.

[7] M. J. Neely, Asynchronous Scheduling for Energy Optimality in Systems with Multiple Servers, *Proc. 46th Annual Conf. on Information Sciences and Systems (CISS)*, March 2012.

[8] M. J. Neely, Online Fractional Programming for Markov Decision Systems, *Proc. Allerton Conf. on Communication, Control, and Computing*, Sept. 2011.

[9] A. Gandhi, S. Doroudi, M. Harchol-Balter, A. Scheller-Wolf. Exact analysis of the M/M/k/setup class of Markov chains via recursive renewal reward, *Proc. ACM Sigmetrics*, pp.153-166, 2013.

[10] D. D. Yao. Dynamic Scheduling via Polymatroid Optimization, *Proceeding Performance Evaluation of Complex Systems: Techniques and Tools*, pp. 89-113, 2002.

[11] C. Li, M. J. Neely, Solving convex optimization with side constraints in a multi-class queue by adaptive $c\mu$ rule. *Queueing System*, vol. 77, no. 3, pp. 331-372, 2014.

[12] A. Gandhi, *Dynamic Server Provisioning for Data Center Power Management.* PhD thesis, 2013.

[13] R. Urgaonkar, U. C. Kozat, K. Igarashi, and M. J. Neely, Dynamic Resource Allocation and Power Management in Virtualized Data Centers, *IEEE/IFIP Network Operations and Management Symposium 2010*, Osaka, Japan, April 2010.

[14] S. Chen, Y. Wang, and M. Pedram. Resource allocation optimization in a data center with energy storage devices. *Proc. 40th Annual Conf. of the IEEE Industrial Electronics Society*, Oct. 2014.

[15] Y. Yao, L. Huang, A. Sharma, L. Golubchik, M. Neely, Data centers power reduction: a two time scale approach for delay tolerant workloads, *Proceedings of INFOCOM*, pp. 1431-1439, 2012.

[16] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, Dynamic right-sizing for power-proportional data centers, *Proceedings IEEE INFOCOM*, 2011.

[17] L. Parolini, B. Sinopoli, and B. Krogh, Reducing Data Center Energy Consumption via Coordinated Cooling and Load Management, *HotPower08: Workshop on Power Aware Computing and Systems ACM*, 2008.

[18] T. Horvath and K. Skadron, Multi-mode energy management for multitier server clusters, *Proceedings of the 17th international conference on Parallel architectures and compilation techniques*, Toronto, Ontario, Canada: ACM, 2008, pp. 270-279.

[19] M. J. Neely, Stability and probability 1 convergence for queueing networks via Lyapunov optimization, *Journal of Applied Mathematics*, doi:10.1155/2012/831909, 2012.

[20] R. Durrett, *Probability: Theory and Examples*, Cambridge University Press, 2013.

[21] B. Fox, Markov renewal programming by linear fractional programming, *SIAM Journal on Applied Mathematics*, vol. 14, no. 6, pp. 1418-1432, Nov. 1966.